arXiv:1404.5358v1 [stat.AP] 22 Apr 2014

A randomized trial in a massive online open course shows people don't know what a statistically significant relationship looks like, but they can learn.

Aaron Fisher, G. Brooke Anderson, Roger Peng, Jeff Leek*

April 23, 2014

Abstract

Scatterplots are the most common way for statisticians, scientists, and the public to visually detect relationships between measured variables. At the same time, and despite widely publicized controversy, P-values remain the most commonly used measure to statistically justify relationships identified between variables. Here we measure the ability to detect statistically significant relationships from scatterplots in a randomized trial of 2,039 students in a statistics massive open online course (MOOC). Each subject was shown a random set of scatterplots and asked to visually determine if the underlying relationships were statistically significant at the $P < 0.05$ level. Subjects correctly classified only 47.4% (95% CI: 45.1%-49.7%) of statistically significant relationships, and 74.6% (95% CI: 72.5%-76.6%) of non-significant relationships. Adding visual aids such as a best fit line or scatterplot smooth increased the probability a relationship was called significant, regardless of whether the relationship was actually significant. Classification of statistically significant relationships improved on repeat attempts of the survey, although classification of non-significant relationships did not. Our results suggest: (1) that evidence-based data analysis can be used to identify weaknesses in theoretical procedures in the hands of average users, (2) data analysts can be trained to improve detection of statistically significant results with practice, but (3) data analysts have incorrect intuition about what statistically significant relationships look like, particularly for small effects. We have built a web tool for people to compare scatterplots with their corresponding p-values which is available here: http://glimmer.rstudio.com/afisher/EDA/.

Over the last two decades there has been a dramatic increase in the amount and variety of data available to scientists, physicians, and business leaders in nearly every area of application. Statistical literacy is now critical for anyone consuming data analytic reports, including scientific papers, newspaper reports [1], legal cases [2], and medical test results [3, 4]. A lack of sufficient training in statistics and data analysis has been responsible for the retraction of high-profile papers, the cancellation of clinical trials, and mistakes in papers used to justify major economic policy initiatives.

Despite the critical importance of statistics and data analysis in modern life, we have relatively little empirical evidence about how statistical tools work in the hands of typical analysts and consumers. The most well-studied statistical tool is the visual display of quantitative information. Previous studies have shown that humans have difficulty interpreting linear measures of correlation [5], are better at judging relative positions than relative angles [6, 7], and view correlations differently when plotted on different scales [5]. These studies show that mathematically correct statistical procedures may have unintended consequences in the hands of users. The real effect of a statistical procedure depends to a large extent on psychology and cognitive function.

Here we perform a large-scale study of the ability of average data analysts to detect statistically significant relationships from scatterplots. Our study compares two of the most common data analysis tasks, making scatterplots and calculating P-values. It has been estimated that as many as 80% of the plots published across all scientific disciplines are scatterplots [8]. At the same time, and despite widely publicized controversy over their use [9], P-values remain the most common choice for reporting a statistical summary of the relationship between two variables in the scientific literature. In the decade 2000-2010, 15,653 P-values were reported in the abstracts of the *The Lancet*, *The Journal of the American Medical Association*, *The New England Journal of Medicine*, *The British Medical Journal*, and *The American Journal of Epidemiology* [10].

*Correspondence to: jleek@jhsph.edu

Data analysts frequently use exploratory scatterplots for model selection and building. Selecting which variables to include in a model can be viewed as visual hypothesis testing where the test statistic is the plot and the measure of significance is human judgement. However, it is not well known how accurately humans can visually classify significance when looking at graphs of raw data. This classification task depends on both understanding what combinations of sample size and effect size constitute significant relationships, and being able to visually distinguish these effect sizes. We performed a set of experiments to (1) estimate the baseline accuracy with which subjects could visually determine if two variables showed a statistically significant relationship; (2) test whether accuracy in visually classifying significance was changed by the number of data points in the plot or way the plot was presented; and (3) test whether accuracy in visually classifying significance improved with practice. Our intuition is that potential improvements with practice would be better explained by an improved cognitive understanding of statistical significance, rather than an improved perceptive ability to distinguish effect sizes.

Our study was conducted as an ungraded component of a statistics massive online open course (MOOC). While MOOCs have previously been used to study MOOCs [11, 12, 13], to our knowledge this is the first example of a MOOC being used to study the practice of science. All students in the class had been exposed to lecture material explaining P-values and null hypothesis significance testing prior to the study. Each participating student was shown a set of bi-variate scatterplots (examples shown in the top panels of Figure 1). The set of plots included eight plots from seven different categories (Table 1), with two plots from the reference category (of which one was significant and one was not) and one plot from each of the other categories (each randomly chosen to be either significant or non-significant). These plot categories (Table 1) were selected to allow analysis of whether students' accuracy in visually classifying significance changed based on the number of data points in the plot, or the plot's presentation style. Each set of plots shown to a user was randomly selected from a library containing 10 plots from each category (code in supplemental material), of which half were statistically significant (P-values from testing the slope coefficient in a linear regression relating X and Y were between .023 and .025; e.g., Figure 1, top panel, left) and half were not statistically significant (P-values between .33 and .35; e.g., Figure 1, top panel, right).

Reference	100 data points (e.g., top panels, Figure 1)
Smaller n	35 data points
Larger n	200 data points
Best-fit line	100 data points, with best fit line added
Lowess	100 data points, with smooth lowess curve added (using R "lowess" function)
Axis Scale	100 data points, with the axis range increased to 1.5 standard deviations outside X and Y variable ranges (e.g., "zoomed out") [5]
Axis Label	100 data points, with fictional X- and Y-axis labels added corresponding to activation in a brain region (e.g., "Cranial Electrode 33 (Standardized)" versus "Cranial Electrode 92 (Standardized)")

Table 1: Plot Categories Shown to Users

For each plot, students were asked to visually classify the significance of the bi-variate relationship shown in the graph (in the examples plots shown in the top panel of Figure 1, the correct answer would have been "statistically significant" for the left plot, for which the P-value of a linear relationship between the X and Y variables is 0.024, and "not statistically significant" for the right plot, for which the P-value is 0.341). All eight plots were shown at the same time and students submitted responses for all plots in a single submission. After submitting their responses, students were shown the correct answers and given the opportunity to retake the survey with a new set of plots. Students were not told any information about the structure of the survey and so were not able to use the structure of the survey (e.g., the fact that one of the "Reference" plots was significant and one was not) to improve their accuracy.

To analyze responses, we created separate models for the probability of correctly visually classifying significance in: (1) graphs that showed two variables with a statistically significant relationship (e.g., Figure 1, top panel, left) and (2) graphs that showed two variables with a statistically non-significant relationship (e.g., Figure 1, top panel, right). These two types of visual classification correspond to the separate accuracy metrics: human sensitivity to significance (accuracy in giving a positive result in cases where a condition is true) and human specificity to non-significance (accuracy in giving a negative result in cases where a condition is false). In this framework, the hypothetical baseline case where humans have no ability to classify significance corresponds to the sensitivity rate being equal to one minus the specificity rate, which means that the probability of visually classifying a plot as significant is unaffected

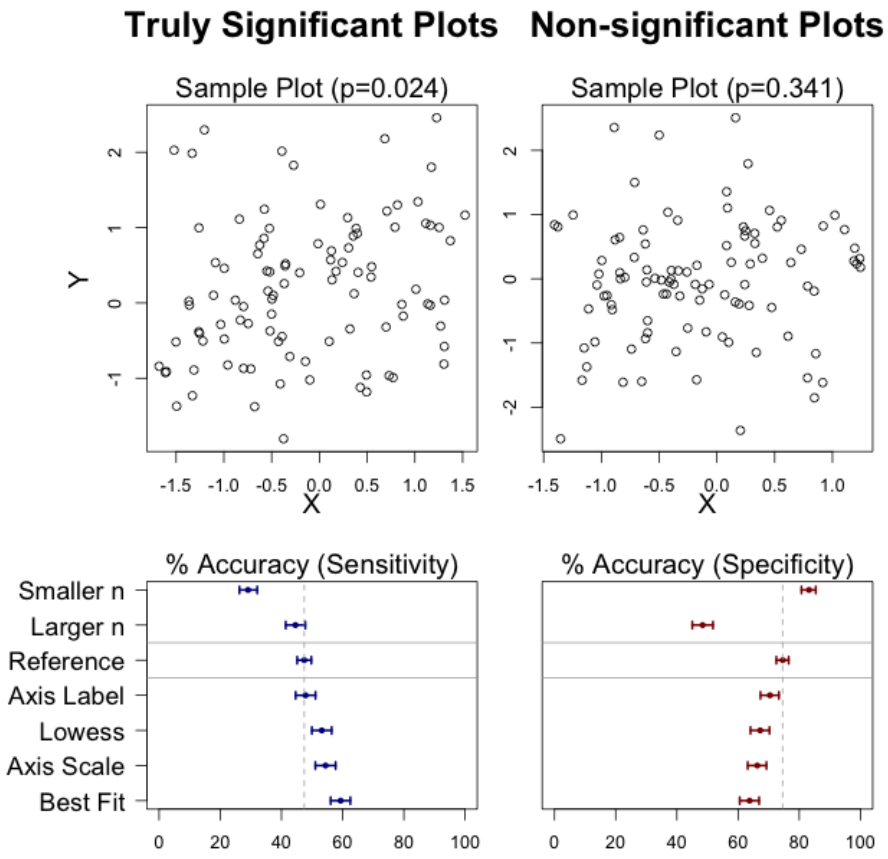

Figure 1: Accuracy of Significance Classifications Under Different Conditions: The top two panels show examples of the plots shown to survey responders. The full library of plots used is provided in the web appendix. For each plot, users were asked to classify the underlying relationship in the data as statistically significant, or not statistically significant. The bottom two panels show point estimates and confidence intervals for classification accuracy under different conditions. Accuracy rates for plots with truly significant underlying relationships are shown on the left (blue), and accuracy rates for plots with non-significant underlying relationships are shown on the right (red).

by the actual significance level of the plot. Accuracy in both metrics was modeled by logistic regressions with person-specific random intercept terms, using the "lme4" package in R (see supplemental materials).

We found that, overall, subjects tended to be conservative in their classifications of significance. In the reference category (100 data; Table 1, examples in Figure 1 top panels), students accurately classified graphs of significant relationships as significant only 47.4% (95% CI: 45.1%-49.7%) of the time (i.e., 47.4% sensitivity) and accurately classified graphs of non-significant relationships as non-significant 74.6% (95% CI: 72.5%-76.6%) of the time (i.e., 74.6% specificity) (Figure 1). Specificity exceeded sensitivity across all of the plot categories presented (Figure 1).

When comparing the reference plot category of 100 datapoints to other plot categories (Table 1), sensitivity and specificity were in some cases significantly changed by the number of points displayed in the graph or the style of graph presentation (Figure 1, bottom panels). Changes to the plots that increased sensitivity correlated with changes that decreased specificity. For example, reducing the number of datapoints shown ("Lower n" plot category) significantly decreased sensitivity (Odds Ratio (OR) = .454, 95%CI: .385-.535) and increased specificity (OR=1.67, 95%CI: 1.39-2.04). Adding visual aids (best-fit line, lowess curve) significantly improved sensitivity (OR=1.62 and 1.26 respectively, with 95%CIs: 1.38-1.89 and 1.08-1.47), but significantly reduced specificity (OR=.600 and .699 respectively, with 95%CIs: .508-.709 and .590-.829). Changing the scale of the axes also increased sensitivity (OR=1.32, 95%CI:

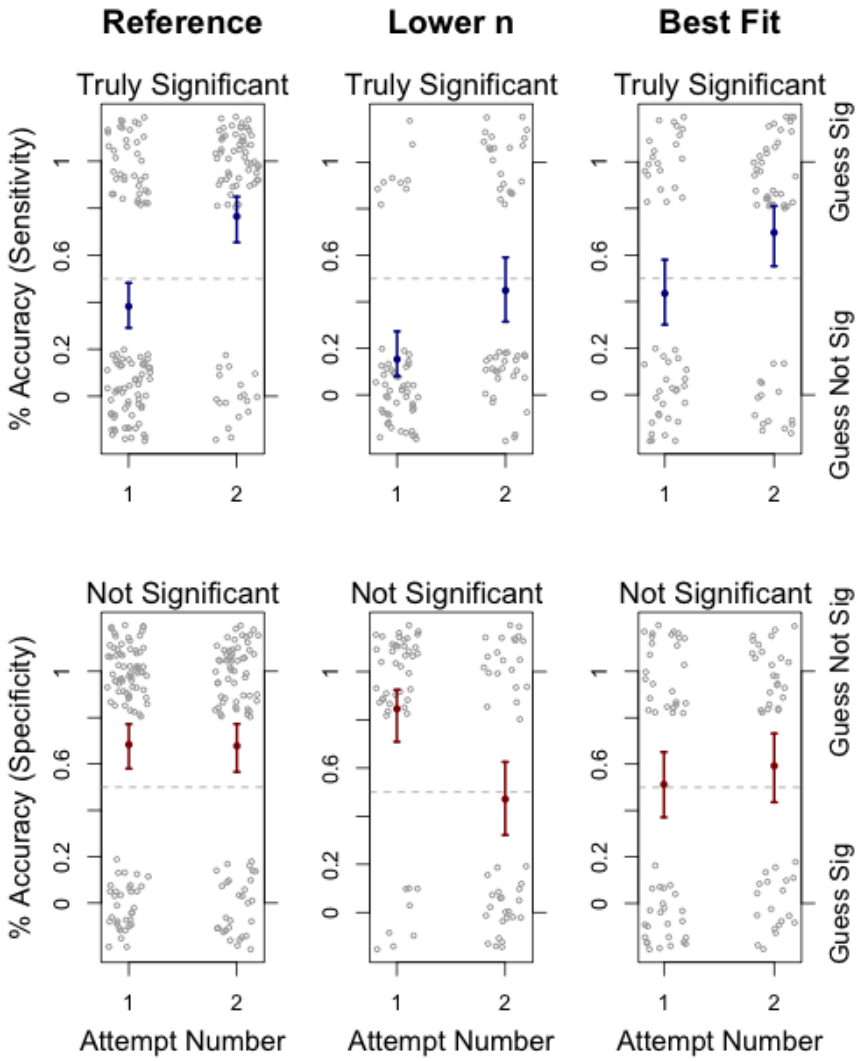

Figure 2: Classification Accuracy on Repeat Attempts of the Survey: Each plot shows point estimates and confidence intervals (in blue and red) for accuracy rates of human visual classifications of statistical significance under different conditions. Plots in the top row correspond to cases where the underlying P-values were significant, while those in the bottom row show cases where the underlying P-values were not significant. Columns of the plots of this figure correspond to different plot presentation styles shown to survey responders (see Table 1). Here, the x axis separates accuracy rates from the first and second attempts of the survey. A jittered version of the raw binary data is shown in gray, corresponding to the secondary y axis labels. For the truly significant underlying P-values, users showed a significant increase in accuracy on the second attempt of the survey for the "Reference," "Lower n," and "Best Fit" presentation styles. For non-significant underlying P-values, accuracy decreased significantly for the "Lower n" category.

1.13-1.55), but decreased specificity (OR=.670, 95%CI: .567-.792). Finally, changing the axes label had no significant effect on sensitivity (OR=1.02, 95%CI: .871-1.19) and only a marginally significant effect on specificity (OR=.811, 95%CI: .682-.965). Because any gain in either specificity or sensitivity tended to come at the cost of the other, none of these plot categories represented a uniform increase in accuracy across all true significance levels of the data underlying the plots.

The exception to this counter-balancing trend came in "Larger n" plots of 200 datapoints, where students showed a significant drop in specificity (OR=.320, 95%CI: .271-.377), and no significant change in sensitivity (OR=.891, 95%CI: .763-1.04). For plots in this category, the probability that users would classify a relationship as significant was fairly similar across truly significant plots and nonsignificant plots. One possible explanation for this is that larger samples require a lower correlation to attain the same significance level. If the correlation becomes imperceptibly small, then the probability that an observer classifies a relationship as significant might be less affected by the true significance level of the plot.

To test if accuracy in visually classifying significance improved with practice, we selected the students who completed the quiz multiple times (101 students) and compared accuracy rates between the students' first and second attempts. We found that sensitivity improved significantly for the "Reference", "Best Fit", and "Lower n" categories (OR=5.27, 2.98, and , 4.51, with 95%CIs: 2.69-10.33, 1.28-6.92, and 1.79-11.37; Figure 2). For the "Reference" and "Best Fit" categories, the sensitivity improvements were not associated with significant changes in specificity, indicating an improvement in overall accuracy in the visual classification of significance. In the "Lower n" plot category however, the increased sensitivity came at the cost of a significant decrease in specificity (OR=.163, 95%CI: .059-.447). For plot in this "Lower n" category, practice did not necessarily improve overall accuracy in visually classifying significance, but rather increased a student's odds of classifying any graph as "significant," regardless of whether the relationship it displayed was truly significant. It is possible that this was due to students over-correcting for their conservatism on their first attempts of the survey. For remaining plot categories ("Higher n", "Axis Label", "Lowess", "Axis Scale"), there were no statistically significant changes in sensitivity or specificity between first and second attempts (Figure 2).

Our research focuses on the question of how accurately statistical significance can be visually perceived in scatterplots of raw data. This work is a logical extension of previous studies on the visual perception of correlation in raw data scatterplots [5, 14, 15],and on the visual perception of plotted confidence intervals in the absence of raw data [16]. The results of this trial are not only relevant towards anyone who wishes to more intuitively understand P-values in scientific literature, but also towards designers and observers of scatterplots. Designers of plots should keep in mind that adding trend lines in a plot tends to make viewers more likely to perceive the underlying relationship as significant, regardless of the relationship's actual significance level, so that they can prevent their plots from misleading viewers. Similarly, viewers of scatterplots may want to slightly discount their perception of statistical significance when trend lines are shown.

Our results also suggest that, on average, readers can improve their ability to visually perceive statistical significance through practice. Our intuition is that this improvement is better explained by an improved understanding of what effect sizes constitute significant relationships, rather than an improved ability to visually distinguish these effect sizes. It would follow that the apparent baseline poor accuracy in visually detecting significance is largely due to a false intuition for what constitutes significant relationships. A broad movement towards practicing the task of visually classifying significance could improve this intuition, and better the efficiency and clarity of communication in science.

This work is also relevant to debate over the misuse of EDA. It has been argued that when EDA and formal hypothesis testing are applied to the same dataset, the "data snooping" committed through EDA process can increase the Type I error rates of the formal hypothesis tests [17]. However, the apparently low sensitivity with which humans can detect statistically significant relationships in scatterplots implies that both the costs of EDA misuse, as well as the benefits of responsibly conducted EDA, may be smaller than expected.

Data analysis involves the application of statistical methods. Our study highlights that even when the theoretical properties of a statistic are well understood, the actual behavior in the hands of data analysts may not be known. Our study highlights the need for placing the practice of data analysis on a firm scientific footing through experimentation. We call this idea evidence based data analysis, as it closely parallels the idea evidence based medicine, the term for scientifically studying the impact of clinical practice. Evidence based data analysis studies the practical efficacy of a broad range of statistical methods when used, sometimes imperfectly, by analysts with different levels of statistical training. Further research in evidence based data analysis may be one way to reduce the well-documented

problems with reproducibility and replicability of complicated data analyses.

To help readers train their sense for P-values, we've created an interactive online application where users can explore the connection between the significance level of a bi-variate relationship and how the data for that relationship appears in a scatterplot. Users can see the visual effect of changing sample size while holding the P-value constant. They can also add lowess curves and best-fit lines to the scatterplot.

`http://glimmer.rstudio.com/afisher/EDA/`

Supplemental Materials

Supplemental materials, including more details on our survey, and code for our analysis, are available at

`https://github.com/aaronjfisher/visual_pvalue/tree/master`

References

[1] Beyth-Marom, R, Fidler, F, & Cumming, G. (2008) Statistical cognition: Towards evidence-based practice in statistics and statistics education. *Statistics Education Research Journal* **7**, 20–39.

[2] Gastwirth, J. L. (1988) *Statistical Reasoning in Law and Public Policy: Volume 2: Tort Law, Evidence and Health.* (Academic Press) Vol. 2.

[3] Schwartz, L. M, Woloshin, S, Black, W. C, & Welch, H. G. (1997) The role of numeracy in understanding the benefit of screening mammography. *Annals of internal medicine* **127**, 966–972.

[4] Sheridan, S. L, Pignone, M. P, & Lewis, C. L. (2003) A randomized comparison of patients' understanding of number needed to treat and other common risk reduction formats. *Journal of General Internal Medicine* **18**, 884–892.

[5] Cleveland, W.S., D. P & McGill, R. (1982) Variables on scatter- plots look more highly correlated when the scales are increased. *Science* **216**.

[6] Heer, J & Bostock, M. (2010) *Crowdsourcing graphical perception: using mechanical turk to assess visualization design.* (ACM), pp. 203–212.

[7] Cleveland, W. S, McGill, R, et al. (1985) Graphical perception and graphical methods for analyzing scientific data. *Science* **229**, 828–833.

[8] Tufte, E. R & Graves-Morris, P. (1983) *The visual display of quantitative information.* (Graphics press Cheshire, CT) Vol. 2.

[9] Nuzzo, R. (2014) Scientific method: Statistical errors. *Nature* **506**, 150–152.

[10] Jager, L. R & Leek, J. T. (2007) Empirical estimates suggest most published medical research is true. *PLoS Med* **4**, e168.

[11] Do, C. B, Chen, Z, Brandman, R, & Koller, D. (2013) *Self-Driven Mastery in Massive Open Online Courses.* (Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA), Vol. 1, pp. 14–16.

[12] Mak, S, Williams, R, & Mackness, J. (2010) *Blogs and forums as communication and learning tools in a MOOC.* (University of Lancaster), pp. 275–285.

[13] Liyanagunawardena, T. R, Adams, A. A, & Williams, S. A. (2013) Moocs: A systematic study of the published literature 2008-2012. *The International Review of Research in Open and Distance Learning* **14**, 202–227.

[14] Meyer, J & Shinar, D. (1992) Estimating correlations from scatterplots. *Human Factors: The Journal of the Human Factors and Ergonomics Society* **34**, 335–349.

[15] Rensink, R. A & Baldridge, G. (2010) *The perception of correlation in scatterplots.* (Wiley Online Library), Vol. 29, pp. 1203–1210.

[16] Belia, S, Fidler, F, Williams, J, & Cumming, G. (2005) Researchers misunderstand confidence intervals and standard error bars. *Psychological methods* **10**, 389.

[17] Berk, R, Brown, L, & Zhao, L. (2010) Statistical inference after model selection. *Journal of Quantitative Criminology* **26**, 217–236.